Preprint
UCRL-JC-154432

Transport Methods Conquering the Seven-Dimensional Mountain

F. R. Graziani, G. L. Olson

This article was submitted to Scales (Science Case for Large Scale Simulation) Washington D. C., June 24-25, 2003

July 21, 2003

U.S. Department of Energy

Lawrence
Livermore
National
Laboratory

This work was performed under the auspices of the United States Department of Energy by the University of California, Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

Transport Methods

Conquering the Seven-Dimensional Mountain

Authors: Frank R. Graziani and Gordon L. Olson

Contributors: Tom Abel, Marv Adams, Eugene Bierly, Fausto Cattaneo, Robert Edwards, Andrew Felmy, Chris Fryer, Brent Gorda, Ivan Hubeny, Stephen Kent, Philip Jardine, Stephen Klippenstein, Willaim Martin, Bronson Messer, Tony Mezzacappa, Jim Morel, John Negele, Peter Nugent, Scott Parker, Scott Studham, Doug Swesty, Robert Tipton, Albert Valocchi

Introduction

In a wide variety of applications, a significant fraction of the momentum and energy present in a physical problem is carried by the transport of particles. Depending on the circumstances, the types of particles might involve some or all of photons, neutrinos, charged particles, or neutrons. In application areas that use transport, the computational time is usually dominated by the transport calculation. Therefore, there is a potential for great synergy; progress in transport algorithms could help quicken the time to solution for many applications.

The complexity, and hence expense, involved in solving the transport problem can be understood by realizing that the general solution to the Boltzmann transport equation is seven dimensional: 3 spatial coordinates, 2 angles, 1 time, and 1 for speed or energy. Low-order approximations to the transport equation are frequently used due in part to physical justification but many times simply because a solution to the full transport problem is too computationally expensive. An example is the diffusion equation, which effectively drops the two angles in phase space by assuming that a linear representation in angle is adequate. Another approximation is the grey approximation, which drops the energy variable by averaging over it. If the grey approximation is applied to the diffusion equation, the expense of solving what amounts to the simplest possible description of transport is roughly equal to the cost of implicit computational fluid dynamics. It is clear therefore, that for those application areas needing some form of transport, fast, accurate and robust transport algorithms can lead to an increase in overall code performance and a decrease in time to solution.

The seven-dimensional nature of transport means that factors of 100 or 1000 improvement in computer speed or memory are quickly absorbed in slightly higher resolution in space, angle, and energy. Therefore, the biggest advances in the last few years and in the next several years will be driven by algorithms. Because transport is an implicit problem requiring iteration, the biggest gains are to be made in finding faster techniques for acceleration to convergence. Some of these acceleration methods are very application specific because they are physics based; others are very general because they address the mathematics of the transport equation. Funding more research in the latter area could have a large impact on many physics applications. Usually it is a collaboration of someone with a tough problem to solve and someone with a new idea that makes the big advances. More heads are needed to continue the progress of the last few years.

Unfortunately, transport as a discipline is not taught in many graduate schools. Students and researchers too often pick up transport theory in pieces on an *ad hoc* basis. Therefore, they don't know the published literature and existing techniques. Knowledge of advances in one application area often takes years to propagate to other application areas.

Applications

Applications abound where transport is required. In astrophysics, the life cycle of the stars, their formation, evolution, and death all require transport. In star formation and evolution for example, the problem is a multi-physics one involving MHD, self-gravity, chemistry, radiation transport, and a host of other phenomena. Supernova core collapse is an example where 3D, multi-group, multi-angle photon and neutrino transport are important in order to model the explosion mechanism. The spectra and light curves generated from a supernova have generated a wealth of data. In order to make a connection between simulation data and observational data and in order to remove systematic errors in supernova standard candle determinations of cosmological parameters, 3D, multi-group, multi-angle radiation transport is required.

The simulation of nuclear reactor science poses a similar set of challenges. In order to move beyond the current state-of-the-art for such calculations, several requirements must be met: 1) a description based on explicit heterogeneous geometry instead of homogenized assemblies; 2) dozens of energy groups instead of two; 3) the use of 3D high-order transport instead of diffusion. These requirements will allow for accurate real-time simulations of new reactor operating characteristics, creating a virtual nuclear reactor test bed. Such a virtual reactor would enable assessments of the impact of new fuel cycles on issues like proliferation and waste repositories. With a 1000-times increase in computer power, accurate virtual reactors could reduce the need to build expensive prototype reactors.

In the broad area of plasma physics, ICF (Inertial Confinement Fusion) and to a lesser extent MFE (Magnetic Fusion Energy) require the accurate modeling of photon and charged particle transport. For ICF, whether one is dealing with direct drive through photon or ion beams or dealing with indirect drive via thermal photons in a hohlraum, the accurate transport of energy around and into tiny capsules requires high-order transport solutions for photons and electrons. For direct drive experiments, simple radiation treatments suffice (i.e. laser ray tracing with multi-group diffusion). Although the radiation treatment can be rather crude, direct drive experiments require sophisticated models of electron transport. In indirect drive such as at NIF, laser energy is converted into thermal x-rays via a hohlraum which in turn is used to drive some target. In order to accurately treat the radiation drive in the hohlraum and its attendant asymmetries will require a radiation transport model with NLTE (Non-Local Thermodynamic Equilibrium) opacities for the hohlraum. The ability to generate NLTE is a tremendous computational challenge. Currently, calculating such opacities in-line comes at a great cost. Typically, the difference between an LTE transport and NLTE transport calculation is a factor of 5. This fact has sparked research into alternatives such as tabulating steady state NLTE opacities or by simplifying the electron population rate equations so that their calculation is fast. However, all of the alternatives suffer from drawbacks which inhibit their widespread use.

In the coming years, simulations of NIF (National Ignition Facility) experiments will be crucial in attaining the goal of ignition. The simulations need to be predictive rather than after-experiment fits; therefore, high-order transport coupled self-consistently to other nonlinear physics is a requirement. With a 1000-fold increase in computer power, these types of simulations are feasible.

There exist a myriad of other applications requiring some form of transport. We close this section by just listing some of them: atmospheric physics, medical diagnostics and treatment, plasma diagnostics, combustion, and non-destructive testing.

Ideas for revolutionary change

As mentioned in the introduction to this chapter, the biggest payoff in algorithm improvements is probably in the area of acceleration of nonlinear iterations to convergence. Pre-conditioners for iterations can be physics-based (diffusion accelerating high-order transport) or mathematics-based (Krylov wrappers around existing iteration techniques). Much progress has been made recently, but much more needs to be accomplished.

Since transport is seven-dimensional, an obvious approach to reducing memory requirements is to use adaptive meshing in all dimensions. In spatial dimensions, much progress has been made with AMR (Adaptive Mesh Refinement) techniques. Currently, most applications use the same spatial mesh for all physics components. In the future, each component may have a mesh optimized for its own needs. Applying these techniques to the angular and energy dimensions of phase space is an area of research that is still in its infancy. The benefits are still to be achieved.

Some of the other important areas of research: Newton-Krylov techniques for tightly coupled self-consistent multi-physics simulations (see the chapter on Multiphysics Techniques), hybrid deterministic and stochastic transport techniques, higher-order time integration, subgrid models for stochastic media simulations, etc.

How would you measure success?

Currently, most of the applications mentioned above do not use full transport but rather some approximation. The approximations currently used are grey or multi-group diffusion and sometimes diffusion with variable Eddington factors. In extreme cases, transport effects are not included at all. These situations arise not because people are ignorant of the importance of transport, but rather its high cost forces code developers to seek cheaper alternatives. Success in this field would be the enabling of high-order transport in applications where it is currently ignored or approximated.

With the removal of severe physics approximations from the codes, a consequence of the above success is the use of virtual simulations as partners with or even replacements for experiments.

Another mark of success is cultural. Success in this area implies an increase in young students doing transport. The ability to attract bright young people to this area is desperately needed if the barriers are to be overcome. Attracting more people to transport can begin by increasing cross fertilization of transport ideas across disciplines. A measure of success on this front would be the increase in communication amongst the atmospheric, astrophysics, and nuclear engineering communities for example.

Barriers

Hardware – Since the solution of the transport problem exists in a multi-dimensional phase space, computer memory is a key issue. In fact historically, the lack of enough memory has pushed algorithms towards operator splitting and the use of layered iterative methods. Other memory issues such as low bandwidth and large latencies and low bandwidth between processors are existing barriers that must be reduced if parallel transport algorithms are to improve in efficiency. Running 3D multi-physics simulations where the transport dominates both the memory and computational cost of performing calculations implies that memory requirements will continue to be an important barrier for the foreseeable future.

Software – Transport theory in isolation has limited value. To be useful, transport solutions must be part of larger multi-physics calculations that are done with huge codes with hundreds of thousands of lines. Frameworks (see related chapter) and component architectures are still in their infancy. There are no reliable, established techniques for building large-scale, multi-physics codes. Software quality engineering (SQE) and assurance (SQA) techniques must become part of every researcher's daily life. Verification of transport components, before and after integration into multi-physics codes, must be taught to all graduate students. Validation of the integrated multi-physics codes by comparing to experiments and data is a shared responsibility of all the contributors to such a code. At a lower level, there are software issues dealing with compilers, unfriendly and non-fault-tolerant parallelization software, etc.

Algorithmic – The biggest algorithmic challenge yet to be discussed is parallelization strategies. Current techniques work with reasonable efficiencies on hundreds of processors, but efficiency degrades significantly on thousands of processors. This is an active area of research and many promising avenues are being pursued. Unfortunately, parallelization algorithms may have to be machine dependent. Each computer has a different mix of communication latencies and bandwidths.

Cultural – Of all the barriers, this class may be the most difficult to solve. What do we mean by cultural barriers? One big barrier is the conflict between a long lead time for code development and the pressure in academia to publish or produce results on a more rapid time scale. The impact is that young faculty who are up for tenure cannot risk getting involved in large code efforts. Also, computational physics at the level needed for large scale computing is under-appreciated in many universities. The consequence, of course, is that without good faculty, students are not trained in these fields. The bottom line is that a future need for good code physicists and engineers experienced in large scale computing may not be met. In order to overcome these barriers, universities must balance the need to "publish or perish" with the long lead times for large scale computing.

Another cultural barrier is isolation. Each field (astrophysics, nuclear reactor engineering, atmospheric physics, etc.) tends to approach transport in its own way with its own specialized terminology. Individuals in a given field communicate with each other but rarely communicate with transport experts outside of their own field. Although there are conferences devoted to transport, they tend to be attended by individuals from a specific field. This culture inhibits communication. In order to overcome this barrier, more funding should be directed to interdisciplinary collaborations. In addition, conferences fostering interdisciplinary communication should be encouraged.

The previous discussion regarding barriers all involved problems that could perhaps be overcome with advances in technology. Reducing cultural barriers on the other hand requires that policies and attitudes at universities and funding agencies be changed. This is perhaps a bigger challenge than the technical barriers.

This work was performed under the auspices of the U.S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.